arXiv:1405.7227v2 [stat.AP] 28 Oct 2014

Bayesian Spatial Change of Support for Count-Valued Survey Data

Jonathan R. Bradley[1], Christopher K. Wikle[2], Scott H. Holan[2]

Abstract

We introduce Bayesian spatial change of support methodology for count-valued survey data with known survey variances. Our proposed methodology is motivated by the American Community Survey (ACS), an ongoing survey administered by the U.S. Census Bureau that provides timely information on several key demographic variables. Specifically, the ACS produces 1-year, 3-year, and 5-year "period-estimates," and corresponding margins of errors, for published demographic and socio-economic variables recorded over predefined geographies within the United States. Despite the availability of these predefined geographies it is often of interest to data users to specify customized user-defined spatial supports. In particular, it is useful to estimate demographic variables defined on "new" spatial supports in "real-time." This problem is known as spatial change of support (COS), which is typically performed under the assumption that the data follows a Gaussian distribution. However, count-valued survey data is naturally non-Gaussian and, hence, we consider modeling these data using a Poisson distribution. Additionally, survey-data are often accompanied by estimates of error, which we incorporate into our analysis. We interpret Poisson count-valued data in small areas as an aggregation of events from a spatial point process. This approach provides us with the flexibility necessary to allow ACS users to consider a variety of spatial supports in "real-time." We demonstrate the effectiveness of our approach through a simulated example as well as through an analysis using public-use ACS data.

Keywords: Aggregation; American Community Survey; Bayesian hierarchical model; Givens angle prior; Markov chain Monte Carlo; Multiscale model; Non-Gaussian.

[1](to whom correspondence should be addressed) Department of Statistics, University of Missouri, 146 Middlebush Hall, Columbia, MO 65211, bradleyjr@missouri.edu

[2]Department of Statistics, University of Missouri, 146 Middlebush Hall, Columbia, MO 65211-6100

1 Introduction

Each year the U.S. Census Bureau samples approximately 3.5 million housing unit addresses through the American Community Survey (ACS). The ACS has replaced the well-known decennial Census long-form, and provides data-users with a more timely data source (Fey, 2005; Bell et al., 2007). One disadvantage of the ACS is that it provides estimates of key demographic variables based on smaller (than the decennial Census) sample sizes. Consequently, ACS estimates often have large margins of error (Speilman et al., 2013).

To improve precision, the ACS produces "period" estimates, which are survey estimates that are aggregated over different time points to reduce the sampling variability (Bell et al., 2007). Specifically, depending on the level of geography, the ACS produces 1-year, 3-year, and 5-year period estimates. Single year period-estimates are provided for geographic regions that consist of populations of 65,000 or larger (e.g., congressional districts and about 800 counties), 3-year period estimates are provided for geographic regions that consist of populations of 20,000 or larger (e.g., all the states, around 1,800 counties, and 900 metropolitan statistical areas), and 5-year period estimates are provided for all population sizes available (e.g., all census tracts, county subdivisions, and counties). Thus, the ACS-defined spatial support, which we call the "source support," becomes finer (in general) as the period increases.

Although ACS period estimates are available on many different source supports, it would be useful to define any desired spatial support. We call a spatial support of interest, that may possibly differ from the source support, a "target support." For example, the Department of City Planning (DCP) in New York City (NYC) use ACS period estimates of demographics, and social characteristics including poverty count-data. They are interested in obtaining estimates of these variables defined on community districts (target support), but instead use aggregate census tracts (source support) (Department of City Planning, 2014), since ACS data are not available on NYC's community districts.

In general, this problem of changing the source support to the target support is known as spatial change of support (COS). Much of the current methodology is based on assuming that the underlying data is Gaussian (see, e.g., Gelfand et al., 2001; Wikle and Berliner, 2005). See Gotway and Young (2002) for a review. Methodologically, we are interested in the case where the data are represented as counts, or the sum of Bernoulli events contained within a given region, which we model using the Poisson distribution. A straightforward methodology to implement in this setting is known as "simple areal interpolation," where values are imputed proportionally based on the areas of each of the target's areal unit (Flowerdew and Green, 1989; Flowerdew et al., 1991; Rogers and Killough, 1991; Flowerdew and Green, 1992, 1994). Although it is easy to implement, measures of uncertainty are not readily available and, hence, it is difficult to ascertain the performance of the estimates defined on the target support or to use the estimated quantities in an inferential setting.

To address this deficiency, Mugglin and Carlin (1998) and Mugglin et al. (1998) use a hierarchical Bayesian approach based on a Poisson data model. Their approach involves partitioning the source support based on the target support, and then estimating parameters defined by each set of this partition. This approach leads to extensive bookkeeping of the partitioning formed by the source support and target support, which makes implementing their Markov chain Monte Carlo (MCMC) procedure difficult. Furthermore, since the parameters of their Bayesian hierarchical model depend on the source support and choice of target support, separate MCMC procedures are required every time a new target support is proposed. Alternatively, we adopt the perspective summarized in Cressie and Wikle (2011, Chap. 4), where count-valued data in small areas are interpreted as an aggregation of events of a spatial point process. This approach allows us to estimate parameters at the point level and aggregate the latent process to any desired target support. This avoids additional MCMC computations if a new target support is introduced, and gives the flexibility required to analyze ACS period-estimates. For a comprehensive comparison between spatial COS using a point process and spatial COS based on partitioning the source and target support into

disjoint regions, see Trevisani and Gelfand (2013).

Although the aforementioned point process aggregation approach is viable for count data recorded without error, there is currently no spatial COS methodology that can be applied to count-valued survey data where estimates of the variance using the survey's sampling design are considered known (e.g., see, Lohr (1999) for more on sampling designs). We use this additional information to account for the variability due to survey sampling within a Bayesian hierarchical modeling (BHM) framework. This is easily done for the Gaussian spatial linear model paradigm (e.g., Fay and Herriot, 1979; Porter et al., 2014) in which, the variance of an additive error term is set equal to the estimated survey variance; hence, this additive error is used to account for variability due to survey sampling. For the non-Gaussian setting this source of variability is typically ignored in the hierarchical modeling framework. Thus, an important component of the proposed BHM is that we account for the variability due to survey sampling. This general problem of spatial COS for count-valued survey data is of independent interest, and the methodology proposed in this paper can be applied to a variety of data sources in addition to ACS.

Survey sampling addresses one source of variability; however, we would also like to include the variability of the latent process and unknown parameters. To take into account the variability of the latent process and unknown parameters we use a BHM, which extends the approach of Wikle and Berliner (2005). The latent process is modeled using a spatial generalized linear mixed model (GLMM) (e.g., see, Besag et al., 1991; Reich et al., 2006; Hughes and Haran, 2013). However, we propose an extension of the spatial random effects covariance parameter model of Hughes and Haran (2013) by incorporating a Givens angle prior (Anderson et al., 1987; Yang and Berger, 1994; Kang and Cressie, 2011) to allow for additional complexity in the spatial covariance. In particular, we introduce a novel Givens angle prior for areal data models.

The remainder of this paper is organized as follows. Section 2, introduces the spatial COS methodology, which includes the proposed BHM. The data model is defined at the point level, which allows for a flexible approach to spatial COS. Additionally, this model is calibrated in a

manner that incorporates known values of the survey variance. In Section 3, we demonstrate the utility of the suggested modeling approach by simulating a stratified random sample from a model that differs from the statistical model used for COS. Here, we show that one can reasonably recover the true unobserved latent mean-count (known from the simulation) using the methodology outlined in Section 2. In Section 4, we then apply the COS methodology to the motivating example of ACS demographic data recorded over NYC. Finally, concluding discussion is provided in Section 5. For convenience of exposition, details surrounding the MCMC algorithm are left to the Appendix.

2 Spatial Change of Support Methodology

Consider data that are recorded at n_ℓ small areas on the ℓ-th spatial scale $\{A_{\ell,i} : i = 1,...,n_\ell\}$, which creates a partitioning of the spatial domain of interest $D \subset \mathbb{R}^d$; that is, $D \equiv \cup_{i=1}^{n_\ell} A_{\ell,i}$ and $A_{\ell,i} \cap A_{\ell,j} = \emptyset$ $(i \neq j)$. Let $n \equiv \sum_\ell n_\ell$, and $\mathbf{Z} \equiv (Z(A_{\ell,i}) : i = 1,...,n_\ell, \ell = 1,...,L)'$ be a generic n-dimensional vector consisting of count-valued survey data $-$ for example, $\mathbf{Z}$ might consist of the number of individuals below the poverty threshold as estimated by ACS.

Notice that it is possible to observe more than one source support; that is, $\ell = 1,...,L$ where the integer $L \geq 1$. In general, the source supports $\ell = 1,...,L$ are decided *a priori* by the statistical agency that releases the public-use survey data. For example, ACS provides data by census tracts and states for some variables and time periods. We order the spatial scales $\ell = 1,...,L$ from the finest resolution (e.g., $\ell = 1$ might correspond to census tracts) to the coarsest resolution (e.g., $\ell = L$ might correspond to states). Furthermore, we assume that these supports may be misaligned (or non-nested), so that $A_{m,j}$ may include all of $A_{\ell,i}$, none of $A_{\ell,i}$, or partially overlap $A_{\ell,i}$, for $i = 1,...,n_\ell$, $j = 1,...,n_m$, and $m \neq \ell$.

The goal of our analysis is to estimate the mean of $Z(\cdot)$ (denoted by $\mu(\cdot)$) evaluated at unobserved geographies $\{B_m : m = 1,...,M\}$ that may possibly be misaligned with the source supports;

in this context, $\{B_m\}$ is called the target support. The target support is problem specific. For example, from Section 1, we see that the DCP is interested in observing ACS estimates by community districts (which we shall denote as $\{B_m : m = 1, ..., M\}$ for this example), despite the fact that ACS does not publish these estimates.

2.1 Spatial Change of Support

We model an element of the generic n-dimensional count-valued data vector $\mathbf{Z}$ by

$$Z(A_{\ell,i})|\delta(\cdot), p(\cdot) \overset{\text{ind}}{\sim} \text{Pois}\left(\int_{A_{\ell,i}} \delta(\mathbf{u})p(\mathbf{u})d\mathbf{u}\right), \ i = 1, ..., n_\ell, \ell = 1, ..., L, \tag{1}$$

where "Pois" stands for Poisson, $\mu(A) \equiv \int_A \delta(\mathbf{u})p(\mathbf{u})d\mathbf{u}$, and the probability that an event is observed at point location $\mathbf{u} \in D$ is $p(\mathbf{u})$. The function $\delta(\cdot)$ represents the unobserved population density in D, which implies that $N_{\ell,i} \equiv \int_{A_{\ell,i}} \delta(\mathbf{u})d\mathbf{u}$, $i = 1, ..., n_\ell$ and $\ell = 1, ..., L$.

The key feature behind the model in (1) is quantified by the term $\int_{A_{\ell,i}} \delta(\mathbf{u})p(\mathbf{u})d\mathbf{u}$. Specifically, the motivating feature of (1) is that count-valued data in a small area A can be interpreted as an aggregation of events of an inhomogeneous spatial point process (Cressie and Wikle, 2011, Chap. 4) with intensity function $\delta(\cdot)p(\cdot)$; in (1), this aggregation is represented by the integral $\int_A \delta(\mathbf{u})p(\mathbf{u})d\mathbf{u}$.

Now, it is important to note that in the applications motivating this work, survey data are not available below the first spatial resolution (geography). Hence, a necessary assumption for this setting is that for $\mathbf{s} \in A_{1,i}$, $p(\mathbf{s}) \equiv p_{1,i}$; $i = 1, ..., n_1$. Consequently, we can write,

$$Z(A_{1,i})|Y_{1,i} \overset{\text{ind}}{\sim} \text{Pois}\left(\exp(Y_{1,i})\right); \ i = 1, ..., n_1, \tag{2}$$

where the distribution of $\{Y_{1,i} \equiv \log(N_{1,i}p_{1,i})\}$, considered the "process" here, shall be specified in Section 2.3.

Similar to (2), one can write the distribution of $Z(A_{\ell,i})|\mu(A_{\ell,j})$ in terms of $\{\mu(A_{1,i})\}$ (or equivalently $\{Y_{1,i}\}$) for $j = 1,...,n_\ell$ and $\ell = 2,...,L$. To see this, notice that

$$\mu(A_{\ell,j}) = \int_{A_{\ell,j}} p(\mathbf{u})\delta(\mathbf{u})d\mathbf{u};\ \ j = 1,...,n_\ell, \ell = 2,...,L,$$

by definition. If we partition $A_{\ell,j}$ into its potential overlap with all $A_{1,i}$ then

$$\begin{aligned}\mu(A_{\ell,j}) &= \sum_{i=1}^{n_1} \int_{A_{\ell,j}\cap A_{1,i}} p(\mathbf{u})\delta(\mathbf{u})d\mathbf{u} = \sum_{i=1}^{n_1} p_{1,i} \int_{A_{\ell,j}\cap A_{1,i}} \delta(\mathbf{u})d\mathbf{u} \\ &= \sum_{i=1}^{n_1} N_{1,i}p_{1,i}\frac{|A_{\ell,j}\cap A_{1,i}|}{|A_{1,i}|} = \sum_{i=1}^{n_1} \mu(A_{\ell,i})\frac{|A_{\ell,j}\cap A_{1,i}|}{|A_{1,i}|};\ \ j = 1,...,n_\ell,\ \ \ell = 2,...,L. \qquad (3)\end{aligned}$$

Let $h_\ell(j,i) \equiv |A_{\ell,j}\cap A_{1,i}|/|A_{1,i}|$, the n_1-dimensional vector $\mathbf{h}_\ell(j) \equiv (h_\ell(j,1),...,h_\ell(j,n_1))'$, the n_1-dimensional random vector $\boldsymbol{\mu} \equiv (\mu(A_{1,i}) : i = 1,...,n_1)'$, and $|A|$ denote the number of elements in the set A. Then, from (3) we have

$$\mu(A_{\ell,j}) = \mathbf{h}_\ell(j)'\boldsymbol{\mu}\,;\ \ j = 1,...,n_\ell, \ell = 2,...,L, \qquad (4)$$

and

$$Z(A_{\ell,j})|\mu(A_{\ell,j}) \overset{\text{ind}}{\sim} \text{Pois}\left(\mathbf{h}_\ell(j)'\boldsymbol{\mu}\right);\ \ j = 1,...,n_\ell, \ell = 2,...,L. \qquad (5)$$

Equation (4) leads naturally to an approach for estimating $\{\mu(B_m)\}$. That is, from (4) we have that

$$\mu(B_m) = \mathbf{h}_B(m)'\boldsymbol{\mu}\,;\ \ m = 1,...,M, \ell = 2,...,L, \qquad (6)$$

where $h_B(m,i) \equiv |B_m\cap A_{1,i}|/|A_{1,i}|$, the n_1-dimensional vector $\mathbf{h}_B(m) \equiv (h_\ell(m,1),...,h_\ell(m,n_1))'$, and recall $\{B_m\}$ is the target support. Thus, the mechanism to change the spatial support from resolutions $\ell = 1,...,L$ to the B-resolution is the vector $\mathbf{h}_B(\cdot)'$ (i.e., "the change of support vector"), which is consistent with the traditional areal interpolation approach discussed in Section

1, but in a context that manages uncertainty. Furthermore, the product form $\mathbf{h}_B(m)'\boldsymbol{\mu}$ allows for straightforward posterior inference (see Section 2.6).

2.2 Incorporating Survey Variances

Since the survey-based variance estimates of var$\{Z(A_{\ell,i})|\mu(A_{\ell,i})\}$, denoted $\sigma^2_{\ell,i}$, are provided by the ACS, we would like to incorporate this information into the statistical model to improve our estimates of $\mu(A_{\ell,i})$. We take a fully Bayesian approach and assume a distribution for each $\sigma^2_{\ell,i}$.

Ideally, from a "data user" perspective the confidential microdata would be available from the statistical agency to jointly model both $\{\sigma^2_{\ell,i}\}$ and $\{Z(A_{\ell,i})\}$. However, the microdata is not made publicly available due to disclosure limitations. Consequently, a joint model for $\{\sigma^2_{\ell,i}\}$ and $\{Z(A_{\ell,i})\}$ can not be specified based on the microdata. See Quick et al. (2014) and the references therein for a comprehensive discussion regarding confidentiality in the spatial setting.

In general, the typical data-user has some knowledge of what to expect from the sampling distribution of $\{\sigma^2_{\ell,i}\}$. For example, it would be reasonable to expect: the sampling distribution of variance estimates to be unimodal and skewed right; the median of the distribution of $\sigma^2_{\ell,i}$ to be (hopefully) close to the true variance of $Z(A_{\ell,i})|\mu(A_{\ell,i})$ (i.e., $\mu(A_{\ell,i})$); and both $\sigma^2_{\ell,i}$ and $Z(A_{\ell,i})$ to be dependent (since they both depend on the same confidential microdata). Hence, we consider the following model for $\sigma^2_{\ell,i}$:

$$\log(\sigma^2_{\ell,i}) = \log(\mu(A_{\ell,i})) + \varepsilon_{\ell,i};\ \ i = 1,...,n_\ell, \ell = 1,...,L, \tag{7}$$

where $\varepsilon_{\ell,i}$ are independent normal random variables with mean zero and unknown variance $\sigma^2_{\varepsilon,\ell,i}$. Equation (7) implies that $\sigma^2_{\ell,i}$ is lognormal (and hence unimodal and skewed right) with median $\mu(A_{\ell,i})$; $i = 1,...,n_\ell$, and $\ell = 1,...,L$. Additionally, Equations (1) and (7) show that both $\{Z(A_{\ell,i})\}$ and $\{\sigma^2_{\ell,i}\}$ are dependent on the latent process $\{\mu(A_{\ell,i})\}$. This is to be expected since $\{\sigma^2_{\ell,i}\}$ and $\{Z(A_{\ell,i})\}$ are dependent on the same confidential microdata, which in turn, are dependent on the

latent process $\mu(A_{\ell,i})$. Finally, (7) also allows our approach to be robust in cases of overdispersion (i.e., the variance is greater than the mean), through the incorporation of an extra random effect. Despite the simplicity of the data model (7), the incorporation of survey variance estimates in this way substantially improved model performance as shown in Section 3.2.

2.3 The Process Model

Similar to Wikle and Berliner (2005), we take a fully Bayesian approach to COS to account for the variability of $\{Y_{1,i}\}$. The unobserved process $\{Y_{1,i} : i = 1, ..., n_\ell\}$ is modeled as

$$Y_{1,i} = \mathbf{x}_i'\boldsymbol{\beta} + \boldsymbol{\psi}_i'\boldsymbol{\eta} + \xi_i;\ i = 1, ..., n_1, \tag{8}$$

where, $\mathbf{x}_i'\boldsymbol{\beta}$ represents "large-scale" variability from the p-dimensional known covariate vector $\mathbf{x}_{\ell,i}$, and where $\boldsymbol{\beta} \in \mathbb{R}^p$ is the associated fixed and unknown p-dimensional parameter vector. The r-dimensional random vector $\boldsymbol{\eta}$ is assumed to be Gaussian with mean-zero and unknown precision matrix $\mathbf{K}$, and $\{\xi_i\}$ are i.i.d. normal with mean zero and unknown variance σ^2_ξ; note that we reserve the word "Gaussian" for random vectors and the word "normal" for random variables.

In principle, the r-dimensional real vector $\boldsymbol{\psi}_i$ can belong to any class of spatial basis functions (e.g., EOFs, wavelets, radial basis functions, etc.); however, we shall use the Moran's I basis functions (see, Griffith, 2000, 2002, 2004; Griffith and Tiefelsdorf, 2007; Hughes and Haran, 2013; Porter et al., 2013). The Moran's I basis functions are a class of functions used to model areal spatial processes in a reduced dimensional space (i.e., $r \ll n$). This feature allows for fast computation of the distribution of $\boldsymbol{\eta}$, which can become computationally expensive for large r (Hughes and Haran, 2013; Porter et al., 2013). This is particularly useful for ACS data, where, for certain spatial supports (e.g., census tracts), the dataset can be quite large (on the order of thousands of areal units).

Now, $\{\boldsymbol{\psi}_i\}$ is defined to be the r eigenvectors of the Moran's I operator (see, Hughes and Haran,

2013) associated with the corresponding r largest eigenvalues. That is, the Moran's I operator is defined as

$$\mathbf{M}(\mathbf{A}) \equiv \left(\mathbf{I}_{n_1} - \mathbf{X}\left(\mathbf{X}'\mathbf{X}\right)^{-1}\mathbf{X}'\right)\mathbf{A}\left(\mathbf{I}_{n_1} - \mathbf{X}\left(\mathbf{X}'\mathbf{X}\right)^{-1}\mathbf{X}'\right),$$

where the $n_1 \times p$ matrix $\mathbf{X} \equiv (\mathbf{x}_i : i = 1,...,n_1)'$, $\mathbf{I}_{n_1}$ is an $n_1 \times n_1$ identity matrix, and $\mathbf{A}$ is the $n_1 \times n_1$ adjacency matrix corresponding to the edges formed by $\{A_{1,i} : i = 1,...,n_1\}$. Let $\mathbf{\Phi}_M \mathbf{\Lambda}_M \mathbf{\Phi}_M'$ be the spectral decomposition of the $n_1 \times n_1$ matrix $\mathbf{M}(\mathbf{A})$. We follow the suggestion of Hughes and Haran (2013) and let r equal (approximately) 10% of the total number of positive eigenvalues of $\mathbf{M}(\mathbf{A})$ (i.e., the positive diagonal elements of $\mathbf{\Lambda}_M$) and the $n \times r$ matrix $\mathbf{\Psi} \equiv (\boldsymbol{\psi}_i : i = 1,...,r)'$ is defined to be the first r columns of $\mathbf{\Phi}_M$. Specifying larger values of r allows one to model higher frequencies of the hidden process (i.e., model higher resolutions of the spatial process), however, specifying a large value of r also increases the time it takes to compute the distribution of $\boldsymbol{\eta}$ (e.g., see, Bradley et al. (2011), for a discussion). In practice, one needs to determine a balance between computational feasibility and model complexity.

2.4 Parameter Models

We propose a novel extension of the spatial random effect covariance parameter model from Hughes and Haran (2013) to allow for increased flexibility. Specifically, we add structure on the prior distribution for $\mathbf{K}$ (i.e., the covariance matrix of the random effect $\boldsymbol{\eta}$). Hughes and Haran (2013) set $\mathbf{K} = \phi \times \mathbf{\Psi}'\mathbf{Q}\mathbf{\Psi}$, where $\phi > 0$. Here, the $n_1 \times r$ matrix $\mathbf{\Psi} \equiv (\boldsymbol{\psi}_i : i = 1,...,n_1)'$, the $n_1 \times n_1$ matrix $\mathbf{Q} \equiv \text{diag}(\mathbf{A1}) - \mathbf{A}$, and $\mathbf{1}$ is a n_ℓ-dimensional vector of ones.

Using the spectral decomposition of $\mathbf{\Psi}'\mathbf{Q}\mathbf{\Psi}$ we have $\mathbf{K} = \mathbf{\Phi}_Q(\phi \times \mathbf{\Lambda}_Q)\mathbf{\Phi}_Q'$, where $\mathbf{\Phi}_Q$ is a $r \times r$ known orthogonal matrix and $\mathbf{\Lambda}_Q$ is a $r \times r$ known diagonal matrix with nonnegative entries. Hence, the model in Hughes and Haran (2013) is based on the assumption that the eigenvalues of $\mathbf{K}$ are known up to a multiplicative constant and the eigenvectors of $\mathbf{K}$ are known. These assumptions are rather strong since only a single parameter is used to describe the parameter space of the $r \times r$

positive-definite matrix $\mathbf{K}$.

Hence, we extend this parameter model by letting $\mathbf{K} = \mathbf{\Phi}(\phi \times \mathbf{\Lambda}_Q)\mathbf{\Phi}'$ and by adding a prior distribution on $\mathbf{\Phi}$. Specifically, we let $\mathbf{\Phi}$ follow a Givens Angle Prior (GAP) (Anderson et al., 1987; Yang and Berger, 1994; Kang and Cressie, 2011). That is, the $r \times r$ matrix $\mathbf{\Phi}$ is decomposed into a Givens rotator product,

$$\mathbf{\Phi} \equiv (\boldsymbol{O}_{1,2} \times \boldsymbol{O}_{1,3} \times \cdots \times \boldsymbol{O}_{1,r}) \times (\boldsymbol{O}_{2,3} \times \cdots \times \boldsymbol{O}_{2,r}) \times \cdots \times \boldsymbol{O}_{r-1,r},$$

where $\boldsymbol{O}_{i,j}$ is a $r \times r$ identity matrix with the (i,i)-th and (j,j)-th element replaced by $\cos(\theta_{i,j})$ and the (i,j)-th ((j,i)-th) element replaced by $-\sin(\theta_{i,j})$ ($\sin(\theta_{i,j})$). Here, $\theta_{i,j} \in [-\pi/2, \pi/2]$, and we place a prior on the logit transformation of the shifted and rescaled $\theta_{i,j}$ denoted as $\zeta_{i,j} \equiv 1/2 + \theta_{i,j}/\pi$. That is,

$$\text{logit}(\zeta_{i,j}) = a + b \times g_{i,j}(\mathbf{\Phi}_Q);\ i < j = 1, ..., r, \tag{9}$$

where $\text{logit}(\zeta_{i,j}) \equiv \log\{\zeta_{i,j}/(1-\zeta_{i,j})\}$, $a, b \in \mathbb{R}$, and $g_{i,j}$ represents the (i,j)-th Givens angle of $\mathbf{\Phi}_Q$ (e.g., see Yang and Berger (1994) for more details on the Givens angle). We then place a vague Gaussian prior on $(a,b)'$. For the remaining parameters we also assume vague prior distributions. Specifically, we assume $\boldsymbol{\beta}$ has a vague Gaussian prior, and ϕ, σ^2_γ, and $\sigma^2_{\varepsilon,\ell,i}$ ($\ell = 1, ..., L$ and $i = 1, ..., n_\ell$) are each assumed to have vague inverse gamma (IG) priors.

Typically, one places independent Gaussian priors on each $\text{logit}(\zeta_{i,j})$ (Anderson et al., 1987; Yang and Berger, 1994; Kang and Cressie, 2011). However, this would lead to a total of $r \times (r-1)/2$ parameters; hence, for moderately large values of r this would lead to estimating many parameters. For example, if $r = 85$ then there are a total of $85 \times 42 = 3{,}570$ parameters that define the eigenvectors of the covariance matrix $\mathbf{K}$. For the parameter model (9), there are only 2 parameters.

2.5 A Summary of the Bayesian Hierarchical Model

The joint distribution of the data, processes, and parameters is written as the product of the following conditional distributions:

$$
\begin{aligned}
&\text{Data Model 1: } Z(A_{\ell,i})|\boldsymbol{\eta},\boldsymbol{\beta},\xi_i \overset{\text{ind}}{\sim} \text{Pois}\{\mu(A_{\ell,i})\};\\
&\text{Data Model 2: } \log(\sigma^2_{\ell,i})|\boldsymbol{\eta},\boldsymbol{\beta},\xi_i \overset{\text{ind}}{\sim} \text{Normal}\{\log(\mu(A_{\ell,i})),\sigma^2_{\varepsilon,\ell,i}\}; \qquad \text{(10a)}\\
&\text{Process Model 1: } \boldsymbol{\eta}|\boldsymbol{\beta},(a,b)',\phi \sim \text{Gaussian}\{\mathbf{0},\mathbf{K}(a,b,\phi)\};\\
&\text{Process Model 2: } \xi_j|\sigma^2_\gamma \sim \text{Normal}\left(0,\sigma^2_\gamma\right);\\
&\text{Parameter Model 1: } \boldsymbol{\beta} \sim \text{Gaussian}\left(\boldsymbol{\mu}_\beta,\sigma^2_\beta\mathbf{I}_p\right);\\
&\text{Parameter Model 2: } \phi \sim \text{IG}(\alpha_\phi,\omega_\phi); \qquad \text{(10b)}\\
&\text{Parameter Model 3: } (a,b)' \sim \text{Gaussian}\{\boldsymbol{\mu}_\Phi,\sigma^2_\Phi\mathbf{I}_2\}; \qquad \text{(10c)}\\
&\text{Parameter Model 4: } \sigma^2_{\varepsilon,\ell,i} \sim \text{IG}(\alpha_\varepsilon,\omega_\varepsilon) \qquad \text{(10d)}\\
&\text{Parameter Model 5: } \sigma^2_\gamma \sim \text{IG}(\alpha_\gamma,\omega_\gamma), \qquad \text{(10e)}
\end{aligned}
$$

where $\ell = 1,...,L, i = 1,...,n_\ell$, and $j = 1,...,n_1$. In practice, to use this algorithm we have to specify values for the hyperparameters $\boldsymbol{\mu}_\beta$, σ^2_β, α_ϕ, ω_ϕ, μ_Φ, σ^2_Φ, α_ε, ω_ε, α_γ, and ω_γ. We choose values for these hyperparameters so that the corresponding prior distribution is flat and, hence, relatively noninformative. Specifically, we let $\boldsymbol{\mu}_\beta = \mathbf{0}$, $\sigma^2_\beta = 10^{15}$, $\alpha_\phi = \omega_\phi = 1$, $\boldsymbol{\mu}_\Phi = (0,1)'$, $\sigma^2_\Phi = 10^{15}$, $\alpha_\varepsilon = \alpha_\gamma = 1$, and $\omega_\varepsilon = \omega_\gamma = 1$. From our experience, this procedure is rather robust to changes in the prior variance values (i.e., changing 10^{15} to 1000); however, performance tends to be sensitive to choices of $\boldsymbol{\mu}_\Phi$. In particular, the procedure performs well when $\boldsymbol{\mu}_\Phi$ is specified close to $(0,1)'$, and appears to do worse as $\boldsymbol{\mu}_\Phi$ is set to values much different than $(0,1)'$. The value $\boldsymbol{\mu}_\Phi = (0,1)'$ corresponds to a standardized CAR model.

Now, using the Metropolis-Hastings algorithm given in the Appendix, we obtain samples from the posterior distributions of $\boldsymbol{\beta}$, $\boldsymbol{\eta}$, and $\{\xi_i\}$, which we denote with $\boldsymbol{\beta}^*_k \equiv (\beta^*_{k,1},...,\beta^*_{k,p})'$, $\boldsymbol{\eta}^*_k \equiv$

$(\eta^*_{k,1},...,\eta^*_{k,r})'$, and $\{\xi_{i,k} : i = 1,...,n_1, k = 1,....,K\}$, respectively; $k = 1,...,K$. In general, we let the superscript "*" denote a sample from the posterior distribution. Posterior samples of $\mu(A_{1,i})$ and $\boldsymbol{\mu}$ are given by $\mu^*_k(A_{1,i}) \equiv \exp\left(\mathbf{x}'_i\boldsymbol{\beta}^*_k + \boldsymbol{\psi}'_i\boldsymbol{\eta}^*_k + \xi_{i,k}\right)$ and $\boldsymbol{\mu}^*_k \equiv (\mu^*_k(A_{1,i}) : i = 1,...,n_1)'$, respectively; $i = 1,...,n_1$ and $k = 1,...,K$. The result in (6) makes it straightforward to produce estimates of $\mu(\cdot)$ defined on the target support $\{B_m\}$; that is, from (6), $\mu^*_k(B_m) = \mathbf{h}_B(m)'\boldsymbol{\mu}^*_k$, where $m = 1,...,M$ and $k = 1,...,K$. Once one obtains the posterior samples $\{\mu^*_k(B_m)\}$, any statistic can be computed from these samples to make inference on $\mu(B_m)$ for $m = 1,...,M$.

There are other types of statistical inference that are of interest besides estimates of $\mu(\cdot)$. For example, in Section 5, we use the posterior predictive p-value (Rubin, 1984; Meng, 1994; Gelman, 1996) to assess the goodness-of-fit of the model. To do this, replicates from the posterior predictive distribution are found by generating K replicates from the distribution of $\mathbf{Z}, \{\sigma^2_{\ell,i}\}|\boldsymbol{\eta}^*_k, \boldsymbol{\beta}^*_k, \boldsymbol{\varepsilon}^*_k$, where $k = 1,...,K$.

Using the K replicates from $\mathbf{Z}, \{\sigma^2_{\ell,i}\}|\boldsymbol{\eta}^*_k, \boldsymbol{\beta}^*_k, \boldsymbol{\varepsilon}^*_k$, denoted $\{\mathbf{Z}^*_k, \{\sigma^2_{\ell,i}\}^*_k : k = 1,...,K\}$, we compute the likelihood ratio to evaluate if the observed dataset represents a typical value from the posterior predictive distribution. Specifically, we compute

$$\frac{[\mathbf{Z}^*_k, \{\sigma^2_{\ell,i}\}^*_k|\boldsymbol{\eta}^*_k, \boldsymbol{\beta}^*_k, \boldsymbol{\varepsilon}^*_k, \phi^*_k, \boldsymbol{\Phi}^*_k]}{[\mathbf{Z}, \{\sigma^2_{\ell,i}\}|\boldsymbol{\eta}^*_k, \boldsymbol{\beta}^*_k, \boldsymbol{\varepsilon}^*_k, \phi^*_k, \boldsymbol{\Phi}^*_k]}; \ k = 1,...,K, \tag{11}$$

where, in general, $[X|Y]$ denotes the p.d.f. of the conditional distribution of X given Y. Then, the posterior predictive p-value is computed as the proportion of K likelihood ratios in (11) that are greater than one. Posterior predictive p-values very close to zero (one) indicate that the model is producing estimates that are not (too) "close" to the data. Hence, posterior predictive p-values near 0.5 are preferable.

3 Simulations

In this section, a simulated example is used to demonstrate that one can accurately estimate the true values of $\{\mu_{\ell,i}\}$ and $\{\mu(A_{\ell+1,j})\}$ using the proposed methodology. Additionally, a simulation study is used to determine: whether it is useful to incorporate estimates of survey variances using (7); whether the Givens angle prior leads to improvements in prediction; and to assess the computational feasibility of the proposed method for spatial COS. To do this, we generate data from a pre-specified survey design that differs from the proposed model. This approach results in simulated data that more closely resemble a survey that one might observe in practice.

We begin by generating a collection of "pseudo-households" denoted by the set $\{\mathbf{u}_j : j = 1,...,N\}$ - by "pseudo-household," we mean that in a specific context $\{\mathbf{u}_j : j = 1,...,N\}$ might represent all the households in some specified geographical region. We generate the set $\{\mathbf{u}_j : j = 1,...,N\}$ by performing the following steps.

1. Divide the square $D \equiv \{\mathbf{u} : \mathbf{u} = (u_1,u_2)', u_1 \in [1,12], u_2 \in [1,12]\} \subset \mathbb{R}^2$ into a 6×6 grid (Figure 1). Then randomly select three grid-cells (using a discrete uniform distribution) to represent "hot spots," which are colored black.

2. For each grid-cell j that is a hot spot (i.e., colored black) sample 5000 points uniformly within a circle with radius two and center equal to the center of the j-th grid-cell. (Note, there are three such grid cells in Figure 1.)

3. For each grid-cell j that is not a hot spot (i.e., colored white) sample 1000 points uniformly within the j-th grid-cell.

Notice that the set $\{\mathbf{u}_j : j = 1,...,N\}$ is generated from a survey-design that differs from the model that we use to fit the data to better match the type of survey data encountered by a typical user that does not have access to microdata.

Each "pseudo-household" in the set $\{\mathbf{u}_j : j = 1,...,N\}$ references a dichotomous variable denoted $\{w(\mathbf{u}_j) : j = 1,...,N\}$. For example, if $\{\mathbf{u}_j : j = 1,...,N\}$ represents each household in the geographic region of interest, then $\{w(\mathbf{u}_j) : j = 1,...,N\}$ consists of zeros and ones, where $w(\mathbf{u}_j) = 1$ ($w(\mathbf{u}_j) = 0$) could indicate that the family that resides in household j falls below (above) the poverty line. To generate the values for $w(\cdot)$, we let the outcome $w(\mathbf{u}) = 1$ have probability 0.5 for each $\mathbf{u} \in \{w(\mathbf{u}_j) : j = 1,...,N\}$.

Suppose a statistical agency would like to collect a stratified random sample of $\{w(\mathbf{u}_j) : j = 1,...,N\}$. In practice, a statistical agency would likely use a more complex survey-design, but a stratified random sample is enough for the purposes of this simulation study. Divide the region D into the collection of areal units $\{A_{1,i} : i = 1,...,90\}$ displayed in Figure 2(a). Here, the areal units $\{A_{1,i} : i = 1,...,90\}$ represent "strata." In practice, it is not always possible to observe $w(\cdot)$ at all N households. Thus, a sample of size 50 is taken from each areal unit in the set $\{A_{1,i}\}$, which we denote as $\{\mathbf{s}_{i,j} : i = 1,...,90, i = 1,...,50\} \subset \{\mathbf{u}_j : j = 1,...,N\}$. In Figure 2(b), we display $\{\mathbf{s}_{i,j}\}$ and indicate where $w(\mathbf{s}) = 1$ (or $= 0$) for each $\mathbf{s} \in \{\mathbf{s}_{i,j}\}$.

Based on the strata in Figure 2(a), the statistical agency then produces estimates of $\mu(A_{\ell,i}) = \sum_{\mathbf{u} \in A_{\ell,i}} w(\mathbf{u})$; notice that from the simulation we know the values of $\mu(\cdot)$. For this simulated example, the statistics used for stratified random samples (Lohr, 1999, pg. 99−102) are computed as:

$$Z(A_{1,i}) = \frac{N_{1,i}}{50} \sum_{\mathbf{s} \in A_{1,i}} w(\mathbf{s});\ \ i = 1,...,90, \tag{12}$$

$$\sigma^2_{1,i} = \left(1 - \frac{50}{N_{1,i}}\right) N^2_{1,i} \frac{\widehat{p}_{1,i}(1 - \widehat{p}_{1,i})}{49};\ \ i = 1,...,90. \tag{13}$$

Here, $\widehat{p}_{1,i} \equiv \sum_{\mathbf{s} \in A_{1,i}} w(\mathbf{s})/50$ is the sample proportion, and $N_{1,i}$ represents the number of elements in the set $\{\mathbf{u}_j : j = 1,...,N\} \cap A_{1,i}$; $i = 1,...,n$. In Figures 2(c) and 2(d), we display $\{Z(A_{\ell,i})\}$ and $\{\sigma^2_{\ell,i}\}$ computed using the realized values of $\{w(\mathbf{s}_{i,j}) : i = 1,...,90, j = 1,...,50\}$ in Figure 2(b) and the strata in Figure 2(a).

3.1 A Simulated Example

Consider a data-user who is interested in obtaining estimates and measures of uncertainty defined on a different support than what is provided by the statistical agency in Figures 2(c)-2(d). Specifically, consider the target support presented in Figure 2(e) which, as seen in Figure 2(f), is misaligned with the source support $\{B_m : m = 1, ..., 36\}$. To use the proposed model to solve this problem we need to define $\{\mathbf{x}_i\}$, r, and prior distribution hyperparameters. Here, we take $\mathbf{x}_i \equiv 1$, we let r equal roughly 10% of the number of positive eigenvalues of the Moran's I operator, and set the prior parameters equal to the choices given in Section 2.6. The MCMC algorithm in the Appendix was run for 15,000 iterations with a burn-in of 5,000 iterations. Convergence of the MCMC algorithm was assessed visually using trace plots of the sample chains, with no lack of convergence detected.

Figure 3(a,c,d,f) displays the posterior mean and posterior variance for $\{\mu(A_{1,i})\}$ and $\{\mu(B_m)\}$, respectively. When comparing the estimates (Figure 3(a, d)) to the truth (Figure 3(b, e)), we see that the posterior mean reproduces the general trend of the true values of $\{\mu(A_{1,i})\}$ and $\{\mu(B_m)\}$. That is, larger values of the posterior mean and larger values of the truth are located near the black grid-cells displayed in Figure 1. This is also seen in Figure 3(g,h), where we plot $\{\mu(A_{\ell,i})\}$ and $\{\mu(B_m)\}$ and the corresponding posterior means. Hence, it appears that the estimates of $\{\mu(A_{\ell,i})\}$ and $\{\mu(B_m)\}$ are close to their respective true values.

3.2 Multiple Simulated Replicates

The BHM in (11) is currently the only stochastic modeling approach for spatial change of support of count-valued survey data (henceforth referred to as "CS" for count-valued survey-data). In contrast, the primary (non-stochastic) modeling approach that one can directly apply to survey data is simple areal interpolation (SI). Thus, we shall compare the CS model to SI.

We also consider comparing to straightforward adjustments to the BHM in (11) to provide

motivation for the data model in (7) and the parameter models in (10b) − (10e). In particular, we assert that using the data model in (7) leads to more precise estimates than simply ignoring the survey variances $\{\sigma^2_{\ell,i}\}$. Hence, we consider an alternative model that removes (10a) from the BHM in (11). We shall refer to this BHM as the variance removed (VR) model. Additionally, a Givens angle prior is proposed to provide greater flexibility for modeling spatial covariances. To determine whether this added model complexity is beneficial we compare to a BHM that uses the Moran's I prior instead. Specifically, we replace the parameter models in (10b) − (10e) with

$$\phi \sim \text{IG}(\alpha_\phi, \gamma_\phi),$$

where $\mathbf{K} = \phi \times \mathbf{\Psi}'\mathbf{Q}\mathbf{\Psi}$. We shall refer to this BHM as the MI model. In what follows, we compare the CS, VR, MI, and SI over 50 replicates of the simulated field $\{Z(A_{\ell,i})\}$.

To evaluate these models over 50 replicates of $\mathbf{Z}$ we define the *paired absolute deviation* (PAD),

$$\text{PAD}(\text{MD}, \mathbf{Z}) \equiv \frac{1}{M}\sum_{m=1}^{M} \text{abs}\left\{\mu(B_m) - E_{\text{MD}}(\mu(B_m)|\mathbf{Z})\right\} - \frac{1}{M}\sum_{m=1}^{M} \text{abs}\left\{\mu(B_m) - E_{\text{CS}}(\mu(B_m)|\mathbf{Z})\right\}, \tag{14}$$

where the function "abs" is the absolute value function, $E_{\text{MD}}(\mu(B_m)|\mathbf{Z})$ is the posterior expectation of $\mu(B_m)$ with respect to model MD, and MD = VR, MI, and SI. Recall that $\{\mu(B_m) : m = 1, ..., 36\}$ in (14) is known from the simulation. If PAD in (14) is positive for a given replicate of the field $\{Z(A_{1,i}) : i = 1, ..., 80\}$, then the CS based estimate of $\{\mu(B_m) : m = 1, ..., 36\}$ is considered "better than" the MD based estimate.

Thus, PAD in (14), can be considered the response in a paired experiment between MD (= VR, MI, and SI) and the CS based estimate of $\{\mu(B_m) : m = 1, ..., 36\}$. Boxplots of PAD in (14) over 50 replicates of $\{Z(A_{1,i}) : i = 1, ..., 80\}$ are given in Figure 4(a). Notice that the boxplots corresponding to MD = VR and MI are almost completely greater than zero indicating that our two modeling decisions (i.e., incorporating survey variances and extending the MI prior) leads to

improvements in spatial change of support. Also, the boxplot corresponding to SI is completely above zero, which indicates that we obtain a marked improvement by using CS over the benchmark SI. These results are further corroborated with p-values from sign-tests, which are 4.53×10^{-14}, 4.53×10^{-14}, and 8.88×10^{-16} for MD = VR, MI, and SI, respectively.

When comparing the three boxplots in Figure 4(a) we also see that MD = VR and MI do not lead to substantial differences between the PAD of SI; specifically, the medians across the different values for MD are similar, and the IQRs are relatively large. Hence, not only do we see that CS leads to improvements over SI, but if we do not incorporate the survey variances and use the extended version of the MI prior, then the estimates of $\{\mu(B_m) : m = 1, ..., 36\}$ do not appear to be substantially different from those generated from SI.

Recall that we use the change of support methodology given in Section 2.1 because it allows for efficient spatial change of support, and it avoids computationally expensive MCMC calculations each time a new target support is introduced. Here, we provide CPU times for the MCMC algorithm and for spatial COS using Equation (6) to support this claim. All computations were carried out in Matlab using a Dell Optiplex 7010 Desktop Computer with a quad-Core 3.40 GHz processor and 8 Gbytes of memory. The CPU time involved with MCMC calculations and spatial COS, for each of the 50 replicates of $\mathbf{Z}$, is summarized in the histograms given in Figure 4 (b,c). These plots illustrate that the MCMC algorithm takes a moderate amount of time (the median is approximately 93.03 seconds), and spatial change of support can be computed in real time; that is, the CPU times are consistently close to the median CPU time given by approximately 0.03 seconds.

4 Application: Poverty in New York City

The Department of City Planning (DCP) in New York City (NYC) is involved in promoting growth in the communities of NYC. They achieve this, in part, by establishing policies and zoning changes for the entire city and for individual neighborhoods. The DCP also provides technical support

related to housing, public transit, public space, and community facilities.

The period estimates from ACS are useful for DCP's mission (Department of City Planning, 2014). In particular, they use ACS period estimates of demographics, and social characteristics. In this example, we consider the 2012 5-year period estimates of the population below the poverty threshold in NYC. These estimates and their corresponding estimated survey variances are displayed in Figure 5. The scatterplot of $\log(\sigma^2_{\ell,i})$ versus $\log(Z(A_{\ell,i}))$ indicates overdispersion. Hence, a Poisson model with random effects appears reasonable for this dataset (e.g., see McCullagh and Nelder, 1989, pg. 198-199, on using random effects to allow for overdispersion in Poisson models). The source support for this example is given by census tracts (Figure 6), and is denoted by $\{A_{1,1},...,A_{1,2166}\}$.

It is of interest to observe the poverty estimates referenced by community districts, which we denote with $\{B_1,...,B_{71}\}$; see Figure 6 for the boundaries of the community districts in NYC. However, ACS period estimates are not available by community districts. In Figure 7, we see that public use microdata areas (PUMA) (aggregate census tracts) are not coterminous with community districts. Despite this misalignment, the DCP treats PUMAs as an approximation to community districts and suggests using PUMAs in place of community districts (Department of City Planning, 2014). This simple approximation can be done using SI, which is presented in Figure 7.

Instead of compromising by using PUMAs to deal with the misaligned spatial supports, we use the proposed spatial COS methodology to change the spatial support of the 2012 5-year period estimates of poverty from census tracts (source support) to community districts (target support). In Figure 8, we provide the posterior means and variances by census tracts and community districts. These predictions are based on the methodology proposed in Section 2 with $\mathbf{x}_{\ell,i} \equiv 1$, $r = 85$, and we set the prior parameters equal to the choices given in Section 2.6. The maximum value of r is 844 and, following the advice of Hughes and Haran (2013), we let r equal roughly 10% of the available Moran's I basis functions yielding $r = 85$. The MCMC algorithm in the Appendix was run for 20,000 iterations with a burn-in of 10,000 iterations. Convergence of the MCMC algorithm

was assessed visually using trace plots of the sample chains, with no lack of convergence detected.

Comparing the model based predictions in Figure 8 to the approximation suggested by DCP presented in Figure 7, we see that the model based estimates of $\{\mu(B_m)\}$ share the same pattern as the DCP approximations. Hence, we appear to be obtaining reasonable estimates of $\{\mu(B_m)\}$ using the model in Section 2. Additionally, notice in Figure 8 that the model based predictions defined on the census tracts "smooth" the observed dataset. That is, the larger values of our predictions correspond to the larger values within the dataset, but not quite as large. We would like to determine if we have oversmoothed (or undersmoothed) the data. To do this, we compute the posterior predictive *p*-value (see Section 2.6). Here, the posterior predictive *p*-value is 0.60, which indicates no lack of fit; that is, we are obtaining a reasonable fit to the data (Rubin, 1984; Meng, 1994; Gelman, 1996) on the source support (i.e., census tracts).

5 Discussion

We have developed a spatial change of support (COS) methodology to analyze count-valued survey data with known survey variances. The COS problem is typically solved under the assumption that the data follows a Gaussian distribution (e.g., Gelfand et al., 2001; Gotway and Young, 2002; Wikle and Berliner, 2005). However, count-valued data is naturally non-Gaussian and, hence, we extend the Gaussian spatial COS methodology to Poisson spatial COS methodology.

There are two key motivating features behind the proposed statistical model. The first motivating feature, is that count-valued data in a small area A can be interpreted as an aggregation of events of an inhomogeneous spatial point process (Cressie and Wikle, 2011, Chap. 4). Hence, parameters are estimated at the point level, which allows one to aggregate the latent process to any desired target support. This avoids implementation of computationally expensive MCMC algorithms every time a target support is introduced, which are required by other spatial COS methodologies that use the Poisson distribution (Mugglin and Carlin, 1998; Mugglin et al., 1998). The second

motivating feature of the proposed approach is that we include a random component to account for the variability due to survey sampling. This allows us to account for much of the variability due to survey sampling, which for the ACS data, can be quite large (Speilman et al., 2013).

In addition to incorporating survey sampling error, we also introduce a BHM that accounts for the variability of the latent process and the parameters. The latent process is modeled using a spatial generalized linear mixed model (GLMM) framework that includes a spatial random effect (e.g., see, Besag et al., 1991; Reich et al., 2006; Hughes and Haran, 2013). We also propose an extension of the spatial random effects covariance parameter model from Hughes and Haran (2013), which allows for additional model complexity in the spatial covariance.

A simulated example is presented in which we simulate data at the point level and use a stratified random sample to produce survey-based estimates. Although the spatial COS methodology differed from the simulation model, it is apparent from the results of the simulation that one can accurately estimate the unobserved ("true") mean using the proposed methodology.

A real data analysis based on demographic data obtained from the American Community Survey (ACS) is used to illustrate the COS methodology. In this example, we observe the estimated population below the poverty level reported at New York City's census tracts. We use the proposed spatial COS methodology to change the spatial support from census tracts to community districts. Then, the posterior predictive p-value is used to evaluate the performance of the estimated mean population below the poverty level. The results of the application suggest that we appear to be obtaining reasonable estimates using the proposed methodology.

The BHM defined in Section 2 uses a reduced rank modeling approach by extending the parameter model from Hughes and Haran (2013). We have chosen to take the suggestion of Hughes and Haran (2013) and use columns that correspond to 10% of the total number of positive eigenvalues. One could make the model more complex and use a stochastic search variable selection (SSVS) algorithm (e.g., see Wikle and Holan, 2011) to choose spatial basis functions. However, the results of our simulation and application showed that this level of model complexity was not necessary.

There are natural avenues for future research building on the proposed methodology. For example, it would be straightforward to include exogenous covariate information within the large-scale variability term of the model. This would allow one to use this BHM for small area estimation problems. Also, there is a natural extension of the spatial random effect $\boldsymbol{\eta}$ to a spatio-temporal random effect $\boldsymbol{\eta}_t$, where t indexes each year within a given time-period. This generalization would involve both a spatial and a temporal change of support and is a subject of future research.

Another possible extension of the application of our BHM would be downscaling; notice there are no restrictions on $\{B_m\}$ to be coarser than $\{A_{1,i}\}$. However, we note that adjustments to our methodology may be necessary, since there is no information available below $\{A_{1,i}\}$. One interesting route would be to incorporate fine-scale information about the process by replacing $\mathbf{Q}$ (in Section 3) with a matrix built from a deterministic model (if available).

Acknowledgments

We would like to thank two anonymous referees, the associate editor, and the editor for their helpful comments. This research was partially supported by the U.S. National Science Foundation (NSF) and the U.S. Census Bureau under NSF grant SES-1132031, funded through the NSF-Census Research Network (NCRN) program.

References

Anderson, T. W., Olkin, I., and Underhill, L. G. (1987). "Generation of random orthogonal matrices." *SIAM Journal on Scientific and Statistical Computing*, 8, 625–629.

Bell, W., Basel, W., Cruse, C., Dalzell, L., Maples, J., OHara, B., and Powers, D. (2007). "Use of ACS data to produce SAIPE model-based estimates of poverty for counties." Tech. rep., U.S. Census Bureau.

Besag, J. E., York, J. C., and Mollič, A. (1991). "Bayesian image restoration, with two applications in spatial statistics (with discussion)." *Annals of the Institute of Statistical Mathematics*, 43, 1–59.

Bradley, J. R., Cressie, N., and Shi, T. (2011). "Selection of rank and basis functions in the Spatial Random Effects model." In *Proceedings of the 2011 Joint Statistical Meetings*, 3393–3406. Alexandria, VA: American Statistical Association.

Cressie, N. and Wikle, C. K. (2011). *Statistics for Spatio-Temporal Data*. Hoboken, NJ: Wiley.

Department of City Planning (2014). "Population: American Community Survey." http://www.nyc.gov/html/dcp/html/census/popacs.shtml. Retrieved March, 2014.

Fay, R. and Herriot, R. (1979). "Estimates of income for small places: an application of James-Stein procedures to census data." *Journal of the American Statistical Association*, 74, 269–277.

Fey, R. E. (2005). "Model-assisted estimation for the American Community Survey." In *Proceedings of the Joint Statistical Meetings*, 3016–3023. Alexandria, VA: American Statistical Association.

Flowerdew, R. and Green, M. (1989). "Data integration: Statistical methods for transferring data between zonal systems." In *Accuracy of Spatial Databases*, eds. M. Goodchild and S. Gopal, 239–248. London: Taylor and Francis.

— (1992). "Developments in areal interpolating methods and GIS." *Annals of Regional Science*, 26, 67–78.

— (1994). "Areal interpolation and types of data." In *Spatial Analysis and GIS*, eds. S. Fotheringham and P. Rogerson, 121–145. London: Taylor and Francis.

Flowerdew, R., Green, M., and Kehris, E. (1991). "Using areal interpolation methods in Geographic information Systems." *Papers in Regional Science*, 70, 303–315.

Gelfand, Zhu, L., and Carlin, B. (2001). "On the change of support problem for spatio-temporal data." *Biostatistics*, 2, 31–45.

Gelman, A. (1996). "Posterior predictive assessment of model fitness via realized discrepancies." *Statistica Sinica*, 6, 733–807.

Gotway and Young (2002). "Combining incompatible spatial data." *Journal of the American Statistical Association*, 97, 632–648.

Griffith, D. (2000). "A linear regression solution to the spatial autocorrelation problem." *Journal of Geographical Systems*, 2, 141–156.

— (2002). "A spatial filtering specification for the auto-Poisson model." *Statistics and Probability Letters*, 58, 245–251.

— (2004). "A spatial filtering specification for the auto-logistic model." *Environment and PLanning A*, 36, 1791–1811.

Griffith, D. and Tiefelsdorf, M. (2007). "Semiparametric filtering of spatial autocorrelation: The eigenvector approach." *Environment and Planning A*, 39, 1193–1221.

Hughes, J. and Haran, M. (2013). "Dimension reduction and alleviation of confounding for spatial generalized linear mixed model." *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 75, 139–159.

Kang, E. and Cressie, N. (2011). "Bayesian inference for the spatial random effects model." *Journal of the American Statistical Association (Theory and Methods)*, 106, 975–983.

Lohr, S. (1999). *Sampling Design and Analysis*. Pacific Grove, CA, USA: Brooks/Cole Publishing Company.

McCullagh, P. and Nelder, J. (1989). *Generalized Linear Models*. London, UK: Chapman and Hall.

Meng, X. (1994). "Posterior predictive p-values." *The Annals of Statistics*, 22, 1142–1160.

Mugglin, A. and Carlin, B. (1998). "Hierarchical modeling in Geographic Information Systems: Population interpolation over incompatible zones." *Journal of Agricultural, Biological, and Environmental Statistics*, 3, 111–130.

Mugglin, A., Carlin, B., Zhu, L., and Conlon, E. (1998). "Bayesian areal interpolation, estimation, and smoothing: An inferential approach for Geographic Information Systems." *Environment and Planning A*, 31, 1337–1352.

Porter, A., Holan, S. H., and Wikle, C. K. (2013). "Small area estimation via multivariate Fay-Herriot models with latent spatial dependence." *arXiv preprint arXiv: 1310.7211*.

Porter, A. T., Holan, S., Wikle, C., and Cressie, N. (2014). "Spatial Fay-Herriot models for small area estimation with functional covariates." *arXiv preprint arXiv: 1303.6668*.

Quick, H., Holan, S., Wikle, C., and Reoter, J. (2014). "Bayesian marked point process modeling for generating fully synthetic public use data with point-referenced geography." *arXiv preprint arXiv: 1407.7795*.

Reich, B., Hodges, J., and Zadnik, V. (2006). "Effects of residual smoothing on the posterior of the fixed effects in disease − mapping models." *Biometrics*, 62, 1197–1206.

Rogers, J. and Killough, G. (1991). "Historical does reconstruction project: Estimating the population at risk." *Health Physics*, 72, 186–194.

Rubin, D. B. (1984). "Bayesianly justifiable and relevant frequency calculations for the applied statistician." *The Annals of Statistics*, 4, 1151–1172.

Speilman, S., Folch, D., and Nagle, N. (2013). "Patterns and causes of uncertainty in the American Community Survey." *Applied Geography*, 46, 147–157.

Trevisani, M. and Gelfand, A. (2013). "Sampling designs and prediction methods for Gaussian spatial processes." In *Advances in Theoretical and Applied Statistics*, eds. N. Torelli, F. Pesarin, and A. Bar-Hen, 269–279. Springer-Verlag Berlin Heidelberg.

Wikle, C. and Berliner, M. (2005). "Combining information across spatial scales." *Technometrics*, 47, 80–91.

Wikle, C. and Holan, S. (2011). "Polynomial nonlinear spatio-temporal integro-difference equation models." *Journal of Time Series Analysis*, 32, 339–350.

Yang, R. and Berger, J. (1994). "Estimation of a covariance matrix using the reference prior." *Annals of Statistics*, 22, 1195–1211.

Appendix: Bayesian inference using Markov chain Monte Carlo

The MCMC procedure can be implemented using a Metropolis-Hastings within Gibbs sampler with updates as follows. The model in (11) yields,

$$
\begin{aligned}
&[\boldsymbol{\beta}, \boldsymbol{\eta}, \{\xi_i\}, \phi, (a,b)', \{\sigma^2_{\varepsilon,\ell,i}\}, \sigma^2_\gamma | \mathbf{Z}, \{\sigma^2_{\ell,i}\}] \\
&\propto \exp\left(\sum_{\ell=1}^{L} \sum_{i=1}^{n_\ell} Z(A_{\ell,i}) Y_{\ell,i} - \exp(Y_{\ell,i}) \right) \times \exp\left(- \sum_{\ell=1}^{L} \sum_{i=1}^{n_\ell} \frac{(\log(\sigma^2_{\ell,i}) - Y_{\ell,i})^2}{2\sigma^2_{\varepsilon,\ell,i}} \right) \\
&\times \exp\left(-\frac{\boldsymbol{\eta}' \mathbf{K}^{-1} \boldsymbol{\eta}}{2} \right) \times \exp\left(- \sum_{i=1}^{n_1} \frac{\xi_i^2}{2\sigma^2_\gamma} \right) \\
&\times \exp\left(-\frac{(\boldsymbol{\beta} - \boldsymbol{\mu}_\beta)'(\boldsymbol{\beta} - \boldsymbol{\mu}_\beta)}{2\sigma^2_\beta} \right) \times \phi^{-\alpha_\phi - 1} \times \exp\left(-\frac{\omega_\phi}{\phi} \right) \\
&\times \exp\left(-\frac{((a,b)' - \boldsymbol{\mu}_\Phi)'((a,b)' - \boldsymbol{\mu}_\Phi)}{2\sigma^2_\Phi} \right) \\
&\times \left\{ \prod_{\ell=1}^{N} \prod_{i=1}^{n_\ell} (\sigma^2_{\varepsilon,\ell,i})^{-\alpha_\varepsilon - 1} \right\} \times \exp\left(-\omega_\varepsilon \sum_{\ell=1}^{L} \sum_{i=1}^{n_\ell} \frac{1}{\sigma^2_{\varepsilon,\ell,i}} \right) \\
&\times (\sigma^2_\gamma)^{-n_1 \alpha_\varepsilon - n_1} \times \exp\left(-\frac{n_1 \omega_\gamma}{\sigma^2_\gamma} \right)
\end{aligned}
\tag{A.1}
$$

Now, organize the processes and parameters into a set $\boldsymbol{\theta} \equiv \{\boldsymbol{\beta}, \boldsymbol{\eta}, \{\xi_i\}, \phi, (a,b)', \{\sigma^2_{\varepsilon,\ell,i}\}, \sigma^2_\gamma\}$. Let $\boldsymbol{\theta}_{-\eta} \equiv \boldsymbol{\theta} \cap \{\boldsymbol{\eta}\}^c$, where "$c$" is the set complement; similarly, one can define $\boldsymbol{\theta}_{-\beta}$, $\boldsymbol{\theta}_{-\xi}$, $\boldsymbol{\theta}_{-\varepsilon,i}$, $\boldsymbol{\theta}_{-(a,b)}$, $\boldsymbol{\theta}_{-\phi}$, and $\boldsymbol{\theta}_{-\gamma}$. Some of the corresponding full-conditionals can not be sampled directly.

Hence, from (A.1) we have

$$
\begin{aligned}
[\boldsymbol{\eta}|\boldsymbol{\theta}_{-\eta},\mathbf{Z},\{\sigma^2_{\ell,i}\}] &\propto \exp\left(\sum_{\ell=1}^{L}\sum_{i=1}^{n_\ell} Z(A_{\ell,i})Y_{\ell,i} - \exp(Y_{\ell,i})\right) \times \exp\left(-\sum_{\ell=1}^{L}\sum_{i=1}^{n_\ell}\frac{(\log(\sigma^2_{\ell,i})-Y_{\ell,i})^2}{2\sigma^2_{\varepsilon,\ell,i}}\right) \\
&\quad \times \exp\left(-\frac{\boldsymbol{\eta}'\mathbf{K}^{-1}\boldsymbol{\eta}}{2}\right), \\
[\boldsymbol{\beta}|\boldsymbol{\theta}_{-\beta},\mathbf{Z},\{\sigma^2_{\ell,i}\}] &\propto \exp\left(\sum_{\ell=1}^{L}\sum_{i=1}^{n_\ell} Z(A_{\ell,i})Y_{\ell,i} - \exp(Y_{\ell,i})\right) \times \exp\left(-\sum_{\ell=1}^{L}\sum_{i=1}^{n_\ell}\frac{(\log(\sigma^2_{\ell,i})-Y_{\ell,i})^2}{2\sigma^2_{\varepsilon,\ell,i}}\right) \\
&\quad \times \exp\left(-\frac{(\boldsymbol{\beta}-\boldsymbol{\mu}_\beta)'(\boldsymbol{\beta}-\boldsymbol{\mu}_\beta)}{2\sigma^2_\beta}\right) \\
[\{\xi_i\}|\boldsymbol{\theta}_{-\xi},\mathbf{Z},\{\sigma^2_{\ell,i}\}] &\propto \exp\left(\sum_{\ell=1}^{L}\sum_{i=1}^{n_\ell} Z(A_{\ell,i})Y_{\ell,i} - \exp(Y_{\ell,i})\right) \times \exp\left(-\sum_{\ell=1}^{L}\sum_{i=1}^{n_\ell}\frac{(\log(\sigma^2_{\ell,i})-Y_{\ell,i})^2}{2\sigma^2_{\varepsilon,\ell,i}}\right) \\
&\quad \times \exp\left(-\sum_{i=1}^{n_1}\frac{\xi_i^2}{2\sigma^2_\gamma}\right) \\
[(a,b)'|\boldsymbol{\theta}_{-(a,b)},\mathbf{Z},\{\sigma^2_{\ell,i}\}] &\propto \exp\left(-\frac{((a,b)'-\boldsymbol{\mu}_\Phi)'((a,b)'-\boldsymbol{\mu}_\Phi)}{2\sigma^2_\Phi}\right) \times \exp\left(-\frac{\boldsymbol{\eta}'\mathbf{K}^{-1}\boldsymbol{\eta}}{2}\right), \\
[\phi|\boldsymbol{\theta}_{-\phi},\mathbf{Z},\{\sigma^2_{\ell,i}\}] &\propto \mathrm{IG}\left(r/2+\alpha_\phi, \omega_\phi + \boldsymbol{\eta}'\boldsymbol{\Psi}'\boldsymbol{\Phi}\boldsymbol{\Lambda}_Q\boldsymbol{\Phi}'\boldsymbol{\Psi}\boldsymbol{\eta}/2\right) \\
[\sigma^2_{\varepsilon,\ell,i}|\boldsymbol{\theta}_{-\varepsilon,\ell,i},\mathbf{Z},\{\sigma^2_{\ell,i}\}] &\propto \mathrm{IG}\left(1/2+\alpha_\varepsilon, \omega_\varepsilon + (\log(\sigma^2_{\ell,i})-Y_{\ell,i})^2/2\right) \\
[\sigma^2_\gamma|\boldsymbol{\theta}_{-\gamma},\mathbf{Z},\{\sigma^2_{\ell,i}\}] &\propto \mathrm{IG}\left(n_1/2+\alpha_\gamma, \omega_\gamma + \sum_{i=1}^{n_1}\xi_i^2/2\right), \qquad \text{(A.2)}
\end{aligned}
$$

which are the full-conditionals needed for the Metropolis-Hastings algorithm. Define

$$
\alpha(\rho_*,\rho^*_{k-1}) \equiv \min\left\{\frac{q_\rho(\rho^*_{k-1})[\rho_*|\boldsymbol{\theta}^*_{-\rho,k-1},\mathbf{Z}]}{q_\rho(\rho_*)[\rho^*_{k-1}|\boldsymbol{\theta}^*_{-\rho,k-1},\mathbf{Z}]}\right\};\ k=1,...,K
$$

where $\rho = \boldsymbol{\eta}, \boldsymbol{\beta}, (a,b)', \phi, \sigma^2_{\varepsilon,\ell,i}$, and σ^2_γ, ρ_* is a value generated from the generic target distribution $q_\rho(\cdot)$, and ρ^*_k (and $\boldsymbol{\theta}^*_{-\rho,k}$) is the k-th MCMC replicate of ρ (and $\boldsymbol{\theta}_{-\rho}$). We now provide step-by-step instructions for running this MCMC algorithm in the list below.

1. Choose initial values for $\boldsymbol{\eta}$, $\boldsymbol{\beta}$, $\{\xi_i\}$, $(a,b)'$, ϕ, $\{\sigma^2_{\varepsilon,\ell,i}\}$, and σ^2_γ and denote them with $\boldsymbol{\eta}^*_0$, $\boldsymbol{\beta}^*_0$, $\{\xi_i\}^*_0$, $(a^*_0,b^*_0)'$, ϕ^*_0, $\{\sigma^2_{\varepsilon,\ell,i}\}^*_0$, and $(\sigma^2_\gamma)^*_0$ respectively. Let $k = 0$.

2. Let $k = k+1$.

3. Update the components of $\boldsymbol{\eta}$.

   (a) Generate a proposed value of $\boldsymbol{\eta}_*$ from target distribution $q_\eta(\boldsymbol{\eta})$.

   (b) Let $\boldsymbol{\eta}_k^* = \boldsymbol{\eta}_*$ with probability $\min\left\{1, \alpha(\boldsymbol{\eta}_*, \boldsymbol{\eta}_{k-1}^*)\right\}$ and let $\boldsymbol{\eta}_k^* = \boldsymbol{\eta}_{k-1}^*$ with probability $1 - \min\left\{1, \alpha(\boldsymbol{\eta}_*, \boldsymbol{\eta}_{k-1}^*)\right\}$.

4. Update the components of $\boldsymbol{\beta}$.

   (a) Generate a proposed value of $\boldsymbol{\beta}_*$ from target distribution $q_\beta(\boldsymbol{\beta})$.

   (b) Let $\boldsymbol{\beta}_k^* = \boldsymbol{\beta}_*$ with probability $\min\left\{1, \alpha(\boldsymbol{\beta}_*, \boldsymbol{\beta}_{k-1}^*)\right\}$ and let $\boldsymbol{\beta}_k^* = \boldsymbol{\beta}_{k-1}^*$ with probability $1 - \min\left\{1, \alpha(\boldsymbol{\beta}_*, \boldsymbol{\beta}_{k-1}^*)\right\}$.

5. Update the components of $\{\xi_i\}$.

   (a) Generate a proposed value of $\{\xi_i\}_*$ from target distribution $q_\beta(\{\xi_i\})$.

   (b) Let $\{\xi_i\}_k^* = \{\xi_i\}_*$ with probability $\min\left\{1, \alpha(\{\xi_i\}_*, \{\xi_i\}_{k-1}^*)\right\}$ and let $\{\xi_i\}_k^* = \{\xi_i\}_{k-1}^*$ with probability $1 - \min\left\{1, \alpha(\{\xi_i\}_*, \{\xi_i\}_{k-1}^*)\right\}$.

6. Update the components of $(a,b)'$.

   (a) Generate a proposed value of $(a_*, b_*)'$ from target distribution $q_{(a,b)}((a,b)')$.

   (b) Let $(a_k^*, b_k^*)' = (a_*, b_*)'$ with probability $\min\left\{1, \alpha((a_*, b_*)', (a_{k-1}^*, b_{k-1}^*)')\right\}$ and let $(a_k^*, b_k^*)' = (a_{k-1}^*, b_{k-1}^*)'$ with probability $1 - \min\left\{1, \alpha((a_*, b_*)', (a_{k-1}^*, b_{k-1}^*)')\right\}$.

7. Update ϕ, $\{\sigma^2_{\varepsilon,\ell,i}\}$, and σ^2_γ.

   (a) Let ϕ_k^* equal to a value generated from the full conditional of ϕ in (A.2).

(b) Generate $\sigma^2_{\varepsilon,\ell,i}$ from it's full conditional in (A.2) for each $\ell = 1,...,L$ and each $i = 1,...,n_\ell$; denote these values as $(\sigma^2_{\varepsilon,\ell,i})^*_k$. Then set $\{\sigma^2_{\varepsilon,\ell,i}\}^*_k = \{(\sigma^2_{\varepsilon,\ell,i})^*_k : \ell = 1,...,L, i = 1,...,n_\ell\}$.

(c) Let $(\sigma^2_\gamma)^*_k$ equal to a value generated from the full conditional of σ^2_γ in (A.2).

8. Repeat steps 2-7 until convergence.

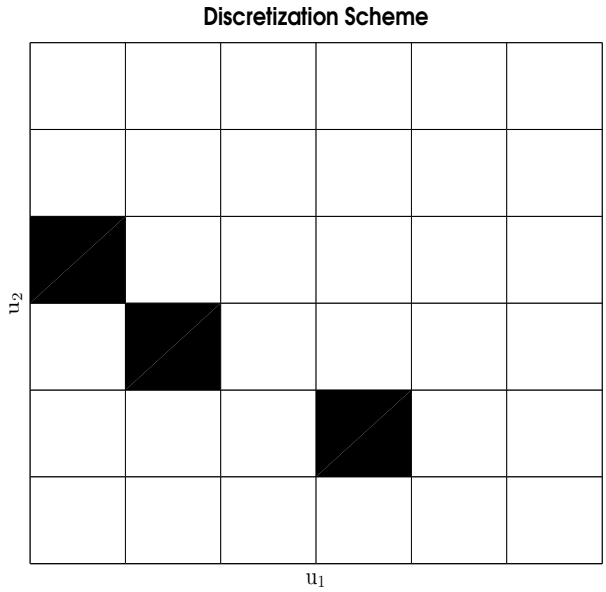

Figure 1: We show a discretization of the spatial domain $D = \{\mathbf{u} : \mathbf{u} = (u_1, u_2)', u_1 \in [1, 12], u_2 \in [1, 12]\}$. Three grid-cells were randomly chosen to represent areas of D that have a high number of psuedo-households (called hot-spots), which are colored back; the remaining cells are colored white.

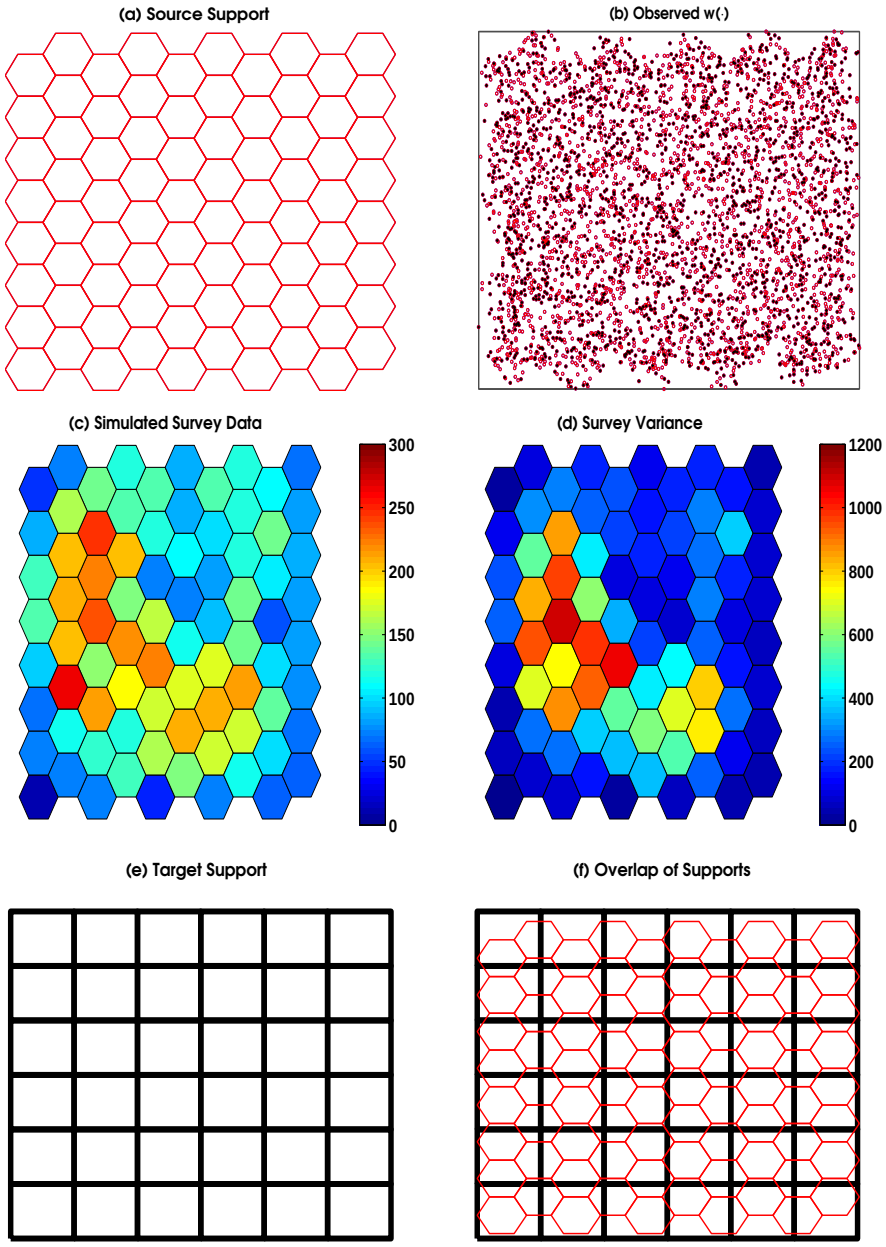

Figure 2: Each polygon in (a) represents an areal unit on the source support (i.e., $A_{\ell,1},...,A_{\ell,90}$). In (b) we sample 50 values of $w(\cdot)$ within each strata to obtain $\{\mathbf{s}_{i,j} : i = 1,...,90, j = 1,...,50\}$ and $\{w(\mathbf{s}_{i,j}) : i = 1,...,90, j = 1,...,50\}$; in (b) we indicate where $w(\mathbf{s}) = 0$ and $w(\mathbf{s}) = 1$ with the colors red and black, respectively. In (c) and (d), we display the sets $\{Z(A_{\ell,i})\}$ and $\{\sigma^2_{\ell,i}\}$ computed using the realized values of $\{w(\mathbf{s}_i) : i = 1,...,n\}$ in (b), Equations (12) and (13), and the strata in (a). Each polygon in (e) represents an areal unit on the target support (i.e., $A_{\ell+1,1},...,A_{\ell+1,36}$). In (f), we have the overlap between the source and target supports.

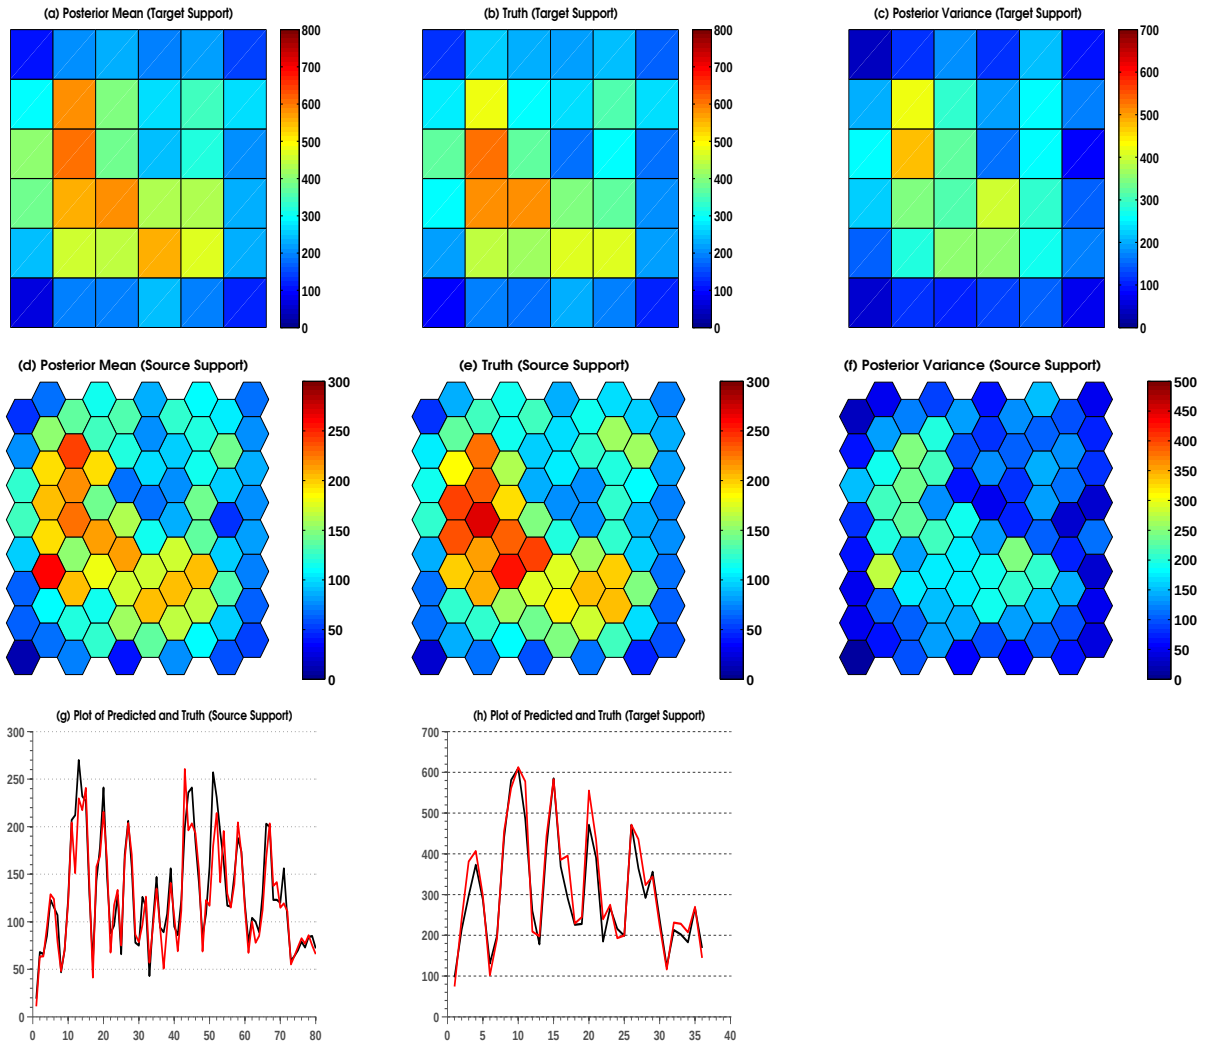
(a) Posterior Mean (Target Support)
(b) Truth (Target Support)
(c) Posterior Variance (Target Support)
(d) Posterior Mean (Source Support)
(e) Truth (Source Support)
(f) Posterior Variance (Source Support)
(g) Plot of Predicted and Truth (Source Support)
(h) Plot of Predicted and Truth (Target Support)

Figure 3: In (a), we present the posterior mean of $\{\mu(A_{\ell,i})\}$, and in (c) we present the posterior variance of $\{\mu(A_{\ell,i})\}$. In (d), we present the posterior mean of $\{\mu(B_m)\}$, and in (f) we present the posterior variance of $\{\mu(B_m)\}$. The respective true fields, known from the simulation, are given in (b) and (e). Note that the color scale between the two rows are different. In (g) we plot $\{\mu(A_{\ell,i})\}$ (black line) and the corresponding posterior mean (red line). In the (h) we plot $\{\mu(B_m)\}$ (black line) and the corresponding posterior mean (red line).

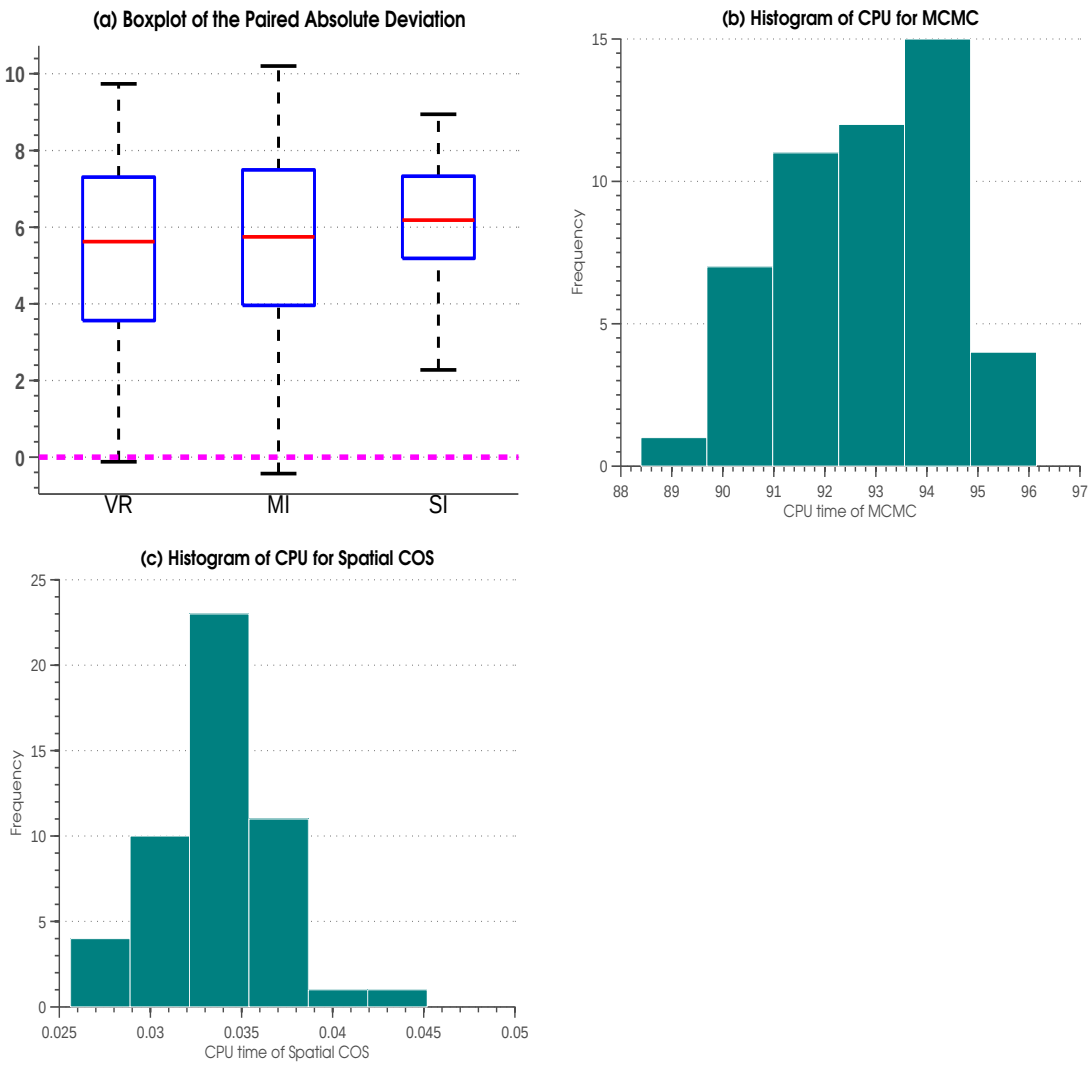

Figure 4: In (a), we present boxplots of PAD in (14), by model MI (Moran's I), VR (variance removed), and SI (simple areal interpolation), and over 50 replicates of **Z**. The steps used to generate each of the 50 **Z**s are outlined in Section 3. The dotted magenta line references a PAD of 0; PAD above (below) this line indicates that model CS (MD) yields better results than MD (CS). In (b) and (c), we present histograms of CPU times (in seconds) over the 50 replicates of **Z**. The steps used to generate each **Z** are outlined in Section 3. In (b), we provide a histogram of the CPU time (in seconds) required for MCMC computation of the BHM in (11). In (c), we provide a histogram of the CPU time (in seconds) required for spatial COS using Equation (6).

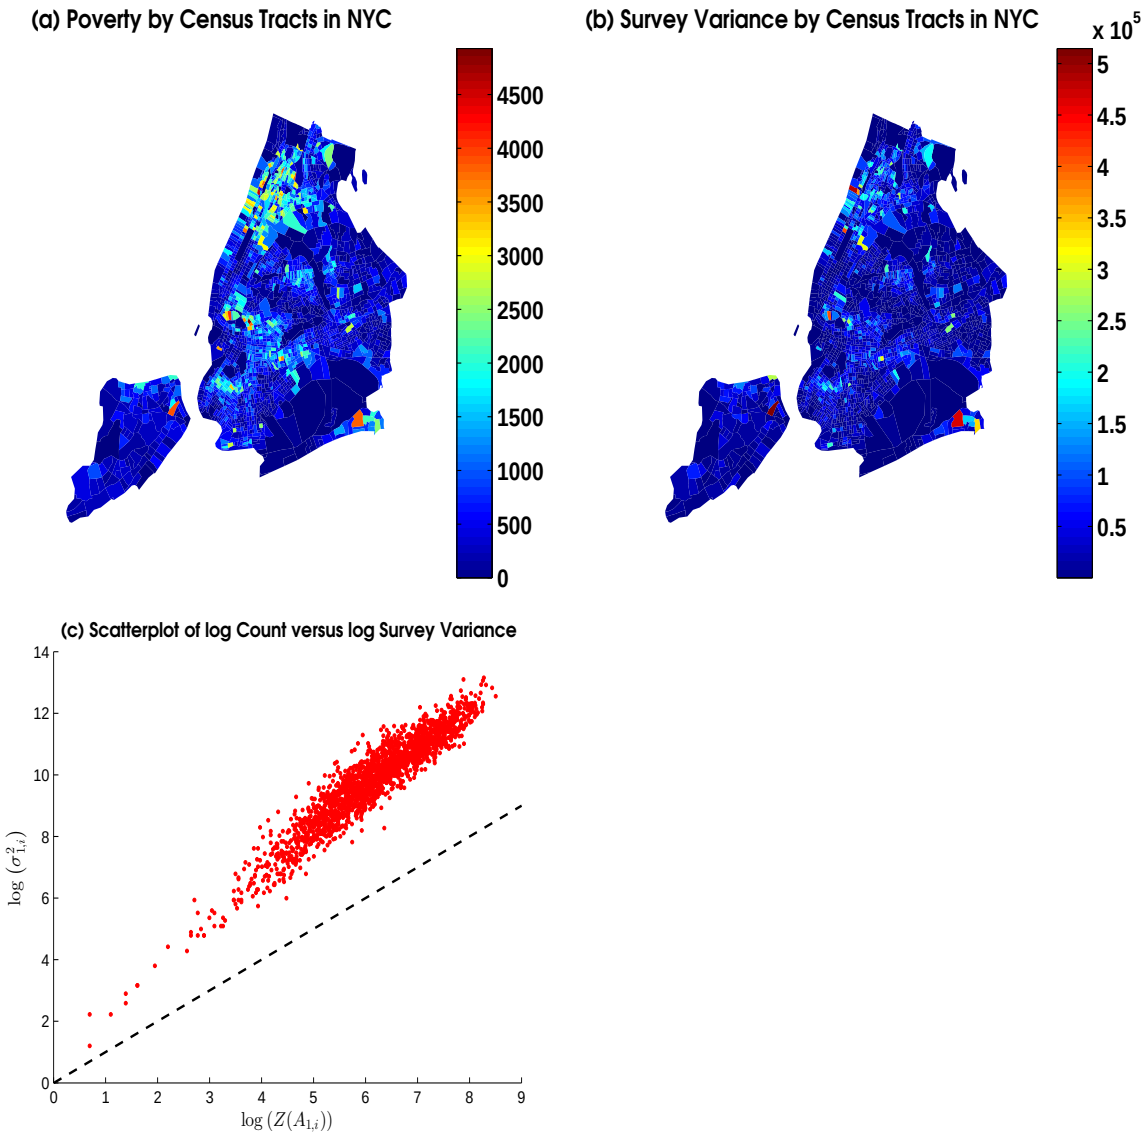

Figure 5: The 2012 ACS 5-year period estimates of the population below the poverty threshold (left column), and the corresponding survey-based variance estimates (right column). These values are recorded over New York City's Census Tracts in (a) and (b). In (c), we provide a scatterplot of the log count data and log survey variance estimate. The dashed black line in (c) represents the line $y = x$.

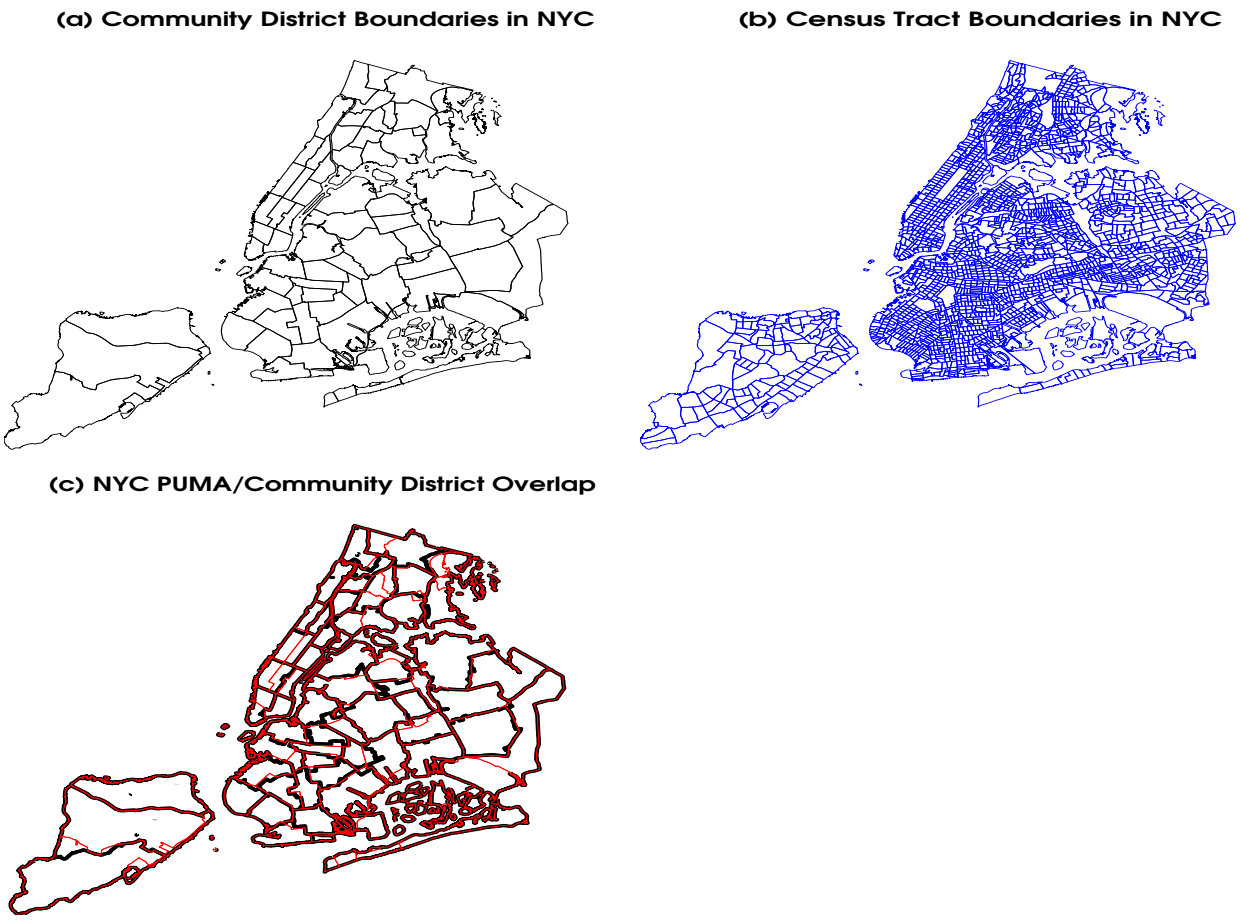

Figure 6: In (a) we display the boundaries of the community districts, which is a target support for this example. In (b) we display the boundaries of the census tracts, which is the source support for this example. In (c), the overlap between PUMA (aggregate census tracts) and community districts are given. The black lines represent PUMA boundaries, and the red lines represent community district boundaries in NYC.

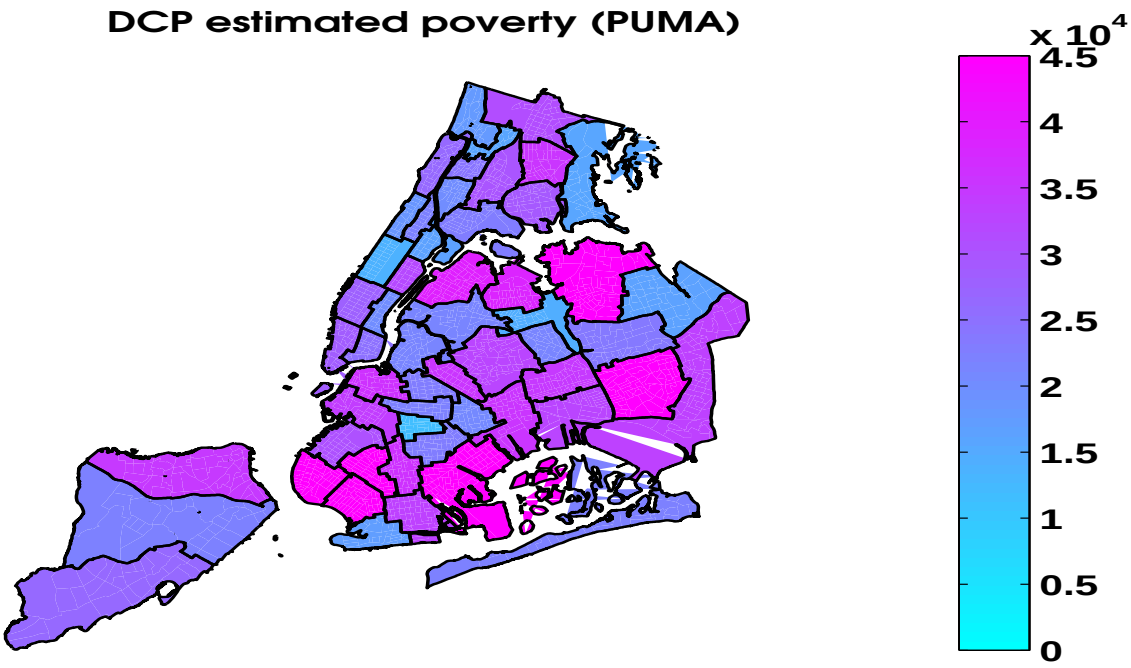

Figure 7: The New York City Department of City Planning approximations of poverty (corresponding to the 2012 5-year period estimate). The values in this map are computed using simple areal interpolation.

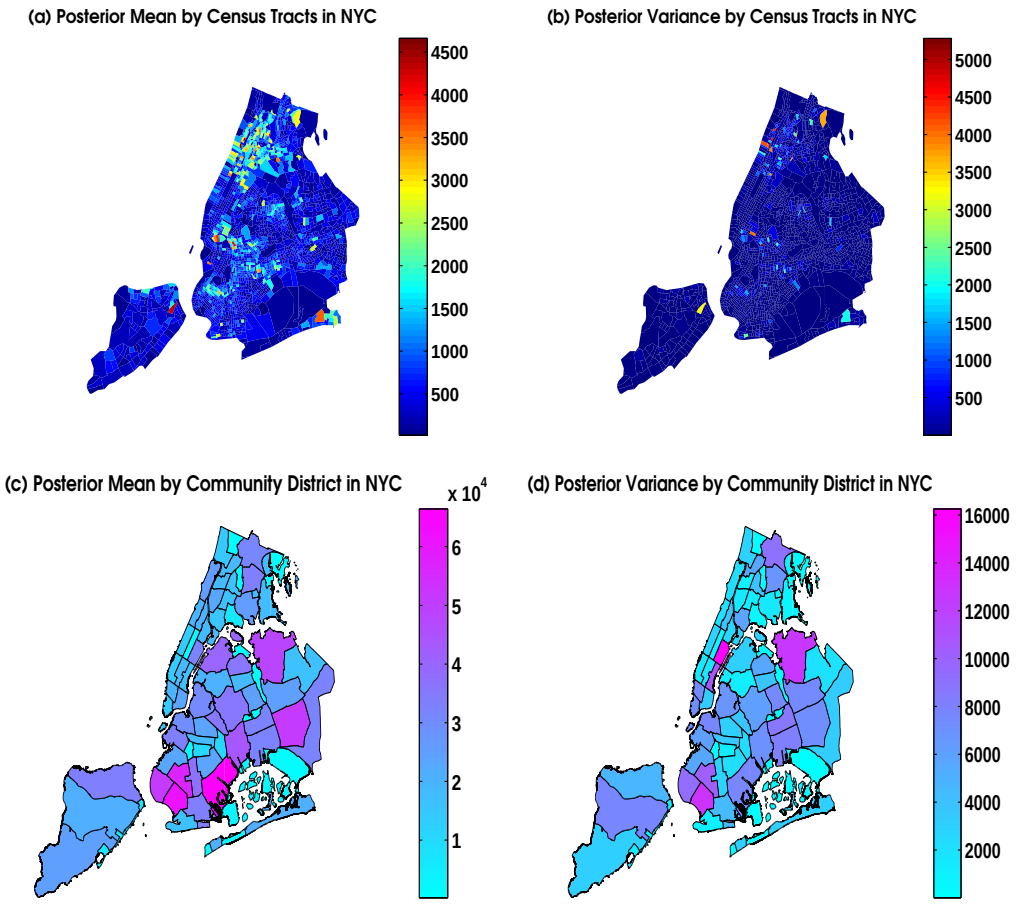

Figure 8: In (a) and (b) we present the posterior means of $\mu(\cdot)$ defined on census tracts (finest source support), and community districts (target support), respectively Each spatial support is indicated by the title headings. In (b) and (d) we present their respective posterior variances. Notice that the color-scales are different for each panel.